

BAYESIAN MODELLING AND QUANTIFICATION OF RAMAN SPECTROSCOPY

BY MATTHEW MOORES^{*}, KIRSTEN GRACIE[†], JAKE CARSON^{*}, KAREN
 FAULDS[†], DUNCAN GRAHAM[†], AND MARK GIROLAMI^{*,‡}


University of Warwick^{}, University of Strathclyde[†] and the Alan Turing
 Institute for Data Science[‡]*

Raman spectroscopy is a technique for detecting and identifying molecules such as DNA. It is sensitive at very low concentrations and can accurately quantify the amount of a given molecule in a sample. The presence of a large, nonuniform background presents a major challenge to analysis of these spectra. We introduce a sequential Monte Carlo (SMC) algorithm to separate the observed spectrum into a series of peaks plus a smoothly-varying baseline, corrupted by additive white noise. Our model-based approach accounts for differences in resolution and experimental conditions. By incorporating this representation into a Bayesian functional regression, we can quantify the relationship between molecular concentration and peak intensity, resulting in an improved estimate of the limit of detection. We also calculate the model evidence using SMC to investigate long-range dependence between peaks.


1. Introduction. Raman spectroscopy can be used to identify molecules by the characteristic scattering of light from a laser. The pattern of peaks in a Raman spectrum corresponds to the vibrational modes of the molecule. The shift in frequency of the photons is proportional to the change in energy state, which is reflected in the locations of the peaks. Surface-enhanced Raman scattering (SERS) is a technique that amplifies the Raman signal using metallic substrates, such as nanoparticles ([Le Ru and Etchegoin, 2009](#)). SERS enables imaging of biological processes at the molecular level, which has many applications in medical diagnostics and therapy ([Butler et al., 2016](#)). For example, a Raman-active dye label can be attached to an antibody targeting a specific protein. The concentration of this protein can then be quantified within a living cell. Fluorescence polymerase chain reaction (PCR) has similar applications, but SERS has superior ability to simultaneously detect multiple labels, with sensitivity in the picomolar range ([Faulds et al., 2008](#); [Gracie et al., 2014](#)). Fig. 1a shows examples of SERS

Primary 92E99; secondary 65D10, 62F15, 62H12

Keywords and phrases: Bayesian model choice, chemometrics, functional data analysis, multivariate calibration, nanotechnology, sequential Monte Carlo



(a) Raw SERS spectra for TAMRA; 5 repeat scans of 5 technical replicates.



(b) Spectral signature for TAMRA, manually baseline-corrected using Grams/AI.

Fig 1: Examples of surface-enhanced Raman scattering (SERS) spectra, before and after baseline correction. Substrate of silver nanoparticles with 100 mW laser, excitation wavelength of 532 nm, 0.01 s accumulation time.

spectra for tetramethylrhodamine (TAMRA), a commercially-available dye.

Each Raman-active dye label has a unique spectral signature, comprised by the locations and amplitudes of the peaks. Fig. 1b shows a spectrum that has been manually baseline-corrected using a multi-point polynomial fit in Grams/AI 7.00 (Thermo Scientific, Waltham, MA). The amplitudes of the peaks are completely dependent on the position of the baseline and vice-versa. Errors in baseline correction can make it difficult to identify and quantify the dyes, particularly when several dyes are present in a multiplexed spectrum (Zhong et al., 2011). Manual baseline correction is a highly painstaking and subjective procedure, therefore automated methods should be preferred.

Many of the existing approaches to automated baseline correction attempt to estimate the location of the baseline independently, without also estimating the peaks. Examples include asymmetric least squares (Boelens, Eilers and Hankemeier, 2005; He et al., 2014), modified polynomial fit (Lieber and Mahadevan-Jansen, 2003), robust local regression (Ruckstuhl et al., 2001), and wavelets (Cai, Zhang and Ben-Amotz, 2001; Galloway, Le Ru and Etchegoin, 2009). See Liland, Almøy and Mevik (2010) for a comparative review.

An iterative algorithm for joint estimation of the baseline and peaks was introduced by De Rooi and Eilers (2012). They combined a penalised spline for the baseline with a mixture model to differentiate between the peaks and the residual noise. The noise was assumed to be Gaussian, while the peaks were modelled using a uniform distribution on the positive real numbers. A drawback of this approach is that it does not provide a continuous model of the peaks, which would account for long-range dependence between neighbouring values in the spectrum and also enable comparison between spectra with different discretisations (x axis). De Rooi and Eilers fit their model using an expectation-maximisation (EM) algorithm (Dempster, Laird and Rubin, 1977), which is known to be highly sensitive to its initial conditions and to be prone to converging to suboptimal local maxima (McLachlan and Krishnan, 2008, ch. 3). These problems can be mitigated using multiple, random starting points combined with annealed likelihood (Zhou and Lange, 2010), but it would be preferable to use an algorithm that explores the full parameter space.

Mixture models have also been used for peak detection and baseline correction in other types of spectroscopy. Razul, Fitzgerald and Andrieu (2003) introduced a reversible-jump Markov chain Monte Carlo (RJ-MCMC) algorithm (Green, 1995) for nuclear emission spectra. The baseline was estimated using a penalised spline, as in De Rooi and Eilers, but the peaks were modelled as squared exponential (Gaussian) kernel functions. The num-

ber and locations of both the knots and the peaks were determined by the trans-dimensional algorithm. They also assumed additive Gaussian noise. Wang et al. (2008) used RJ-MCMC to fit a similar model to mass spectrometry. Kim et al. (2014) used a mixture model for detection of two-dimensional peaks in gas chromatography.

Raman spectra arise due to inelastic scattering of light, therefore the squared exponential kernel functions that were used by Razul, Fitzgerald and Andrieu (2003) and Wang et al. (2008) might not be a suitable model for this data. Previous studies, such as Goldberg and Pershan (1973); Spanier et al. (2001), have concluded that the intrinsic line shape of Raman peaks follows a Lorentzian (Cauchy) function. Both of these functions offer a physical interpretation. Doppler broadening is a result of the emitted photons being red (blue) shifted due to particles moving away from (towards) the sensor. Since the particles are undergoing Brownian motion, this gives rise to a squared exponential function. Collisional broadening occurs due to collisions between particles, which effectively lower the characteristic time of the emission process. As a result of the uncertainty principle this increases the uncertainty in the energy of the emitted photons, which is described by a Lorentzian function. The heavier tails of the Lorentzian would imply long-range dependence between peaks. Failure to account for this would introduce bias, particularly if quantification was based on a single peak in isolation.

In this paper, we consider two alternative functions to model the broadening of the peaks, either squared exponential or Lorentzian. Using Bayesian model choice, we determine which of these functions is better supported by the observed data. We perform a simulation study to demonstrate the effect of model misspecification on parameter estimation. We also need to account for the other components of SERS spectra, the background and residual noise. As in previous methods, we represent the baseline as a penalised cubic spline and assume additive white noise. We extend this model to quantify the relationship between peak intensity and dye concentration for calibration studies, where the concentration levels are systematically varied. We fit this model using a sequential Monte Carlo (SMC) algorithm, which is robust to suboptimal local maxima in the multimodal posterior distribution (Chopin, 2002; Del Moral, Doucet and Jasra, 2006). We have implemented this algorithm as an R package, which is available in Supplement A.

The remainder of this article is organised as follows. We describe the functional model and our informative priors in Section 2. MCMC and SMC algorithms for fitting this model and estimating the model evidence are described in Section 3. These methods are demonstrated using SERS datasets in Section 4. We conclude our article with a discussion.

2. Functional Model. A SERS spectrum is discretised into a multi-variate observation that is highly collinear, hence it lends itself to a reduced-rank representation. Our approach is a form of functional data analysis (Ramsay and Silverman, 2005), where the observed signal is represented using continuous functions. We decompose the spectrum into three major components:

$$(2.1) \quad y_i(\tilde{\nu}) = \xi_i(\tilde{\nu}) + s(\tilde{\nu}) + \epsilon$$

where $y_i(\tilde{\nu})$ is a hyperspectral observation that has been discretised at a number of light frequencies or wavenumbers, $\nu_j \in \tilde{\nu}$. Multiple observations are represented as a matrix $\mathbf{Y}_{1:n_y, \tilde{\nu}}$. The baseline $\xi_i(\tilde{\nu})$ is a smoothly-varying, continuous function that is mainly due to background fluorescence. We assume that ϵ is zero mean, additive white noise with constant variance:

$$(2.2) \quad \epsilon_{i,j} \sim \mathcal{N}(0, \sigma_\epsilon^2)$$

This assumption could be relaxed by allowing for autocorrelated residuals, as in Chib (1993). The Raman peaks are represented as an additive mixture of radial basis functions:

$$(2.3) \quad s(\tilde{\nu}) = \sum_{p=1}^P f(\tilde{\nu} \mid \ell_p, A_p, \varphi_p)$$

where ℓ_p is the location of peak p , A_p is its amplitude, and φ_p is a scale parameter that controls the width or broadening of the peak.

One possible choice to represent the shape of the peaks would be a squared exponential function:

$$(2.4) \quad f(\nu_j \mid \ell_p, A_p, \varphi_p) = A_p \exp \left\{ -\frac{(\nu_j - \ell_p)^2}{2\varphi_p^2} \right\}$$

This function was used by Razul, Fitzgerald and Andrieu (2003) for nuclear spectroscopy and also by Wang et al. (2008) for mass spectroscopy. Under this model, the full width at half maximum (FWHM) of each peak can be estimated as $2\sqrt{2 \ln 2} \varphi_p$. Previous studies (Goldberg and Pershan, 1973; Spanier et al., 2001; Lieber and Mahadevan-Jansen, 2003) suggest that Lorentzian functions would be a better model for Raman spectroscopy:

$$(2.5) \quad f(\nu_j \mid \ell_p, A_p, \varphi_p) = A_p \frac{\varphi_p^2}{(\nu_j - \ell_p)^2 + \varphi_p^2}$$

Under this model, the FWHM is given by $2\varphi_p$. The Lorentzian function decays slower than the squared exponential, which would imply long-range dependence in the spectral signature. In Section 3.2, we describe how Bayesian model choice can be used to determine which of these alternatives is better supported by the observed data.

The amplitudes of the peaks are linearly related to the concentration of the dye, from the limit of detection (LOD) up to monolayer coverage of the nanoparticle surface (MLC). This relationship can be expressed as a regression model:

$$(2.6) \quad A_p = \beta_p c_i, \quad c_{LOD} \leq c_i < c_{MLC}$$


where β_p is a regression coefficient and c_i is the nanomolar (nM) concentration of the dye in the i th observation. Monolayer coverage refers to the saturation point, thus c_{MLC} is dependent on the size and shape of the nanoparticles, as well as the orientation of the dye on the nanoparticle surface. Jones et al. (1999) calculated that a final concentration of 10 nM is required for monolayer coverage of silver nanoparticles with average diameter of 27 nm. The spectra that we analyse in Sect. 4.2 were obtained from much larger nanoparticles, with average diameter of 78 nm, therefore we would expect c_{MLC} to be higher for this dataset. We do not observe any evidence of saturation in our spectra, even for concentrations up to 24.7 nM.

This alternative formulation of the model can be used to analyse LOD or dilution studies, where SERS spectra are obtained for a range of known values of c_i . The limit of detection is usually estimated using a minimum signal-to-noise ratio (SNR) of 3:


$$(2.7) \quad c_{LOD} = \frac{3\sigma_\epsilon}{\beta_p}$$

The LOD is specific to each peak, thus the number of peaks P that are visible in a SERS spectrum is dependent on the dye concentration. Typically only the largest peak is used to estimate the LOD, but in Sect. 4.2 we provide highest posterior density (HPD) intervals for c_{LOD} jointly for every peak in the spectral signature. The posterior distribution for β_p can also be used to scale the prior for A_p , or to estimate c_i in independent observations of the same dye. In chemometrics this is known as multivariate calibration (Varmuza and Filzmoser, 2009).


2.1. *Priors.* The model described by (2.1) is not identifiable, without some means of distinguishing between the components of the signal. The




(a) Lognormal prior for scales of squared exponential peaks.



(b) Truncated normal prior for amplitudes of squared exponential peaks.



(c) Lognormal prior for scales of Lorentzian peaks.



(d) Truncated normal prior for amplitudes of Lorentzian peaks.

Fig 2: Informative priors for the parameters of the SERS peaks, derived from manual baseline correction and peak fitting of Cy3, TAMRA and FAM spectra using Grams/AI 7.00.

only way to achieve strong identifiability would be to estimate the peaks and baseline separately, while holding the estimate of one component fixed. Methods that perform baseline subtraction, noise correction and peak estimation as separate steps can often result in suboptimal solutions. This is particularly the case when a single estimate is provided, without any measure of uncertainty. Errors that are introduced in one stage of the process are propagated forward, without any possibility of correction. Instead, we include all of the components of the signal in a joint model. Our model incorporates prior information to improve identifiability and regularise the solution.

The hyperparameters for these priors were estimated by manual baseline correction and peak fitting in Grams/AI, the same procedure that was used to produce Fig. 1b. We selected three representative spectra, one each of TAMRA, fluorescein (FAM), and cyanine (Cy3), from an independent set of experimental data that had been previously analysed by Gracie et al. (2014). We fitted both squared exponential and Lorentzian peaks to obtain distributions for the amplitude and scale parameters. All of these dyes are well studied, so the number P and locations ℓ_p of the peaks were assumed known. Similar methods could also be used to derive informative priors for these parameters, to perform peak detection as well as quantification. The RJ-MCMC algorithms of Razul, Fitzgerald and Andrieu (2003) and Wang et al. (2008) could be used to fit such a model.

We selected a lognormal distribution for the prior on the scale parameter, since this provided a good fit to the peaks in our training data. The median of the scales was 16.47 for squared exponential peaks and the standard deviation of $\log\{\varphi_p\}$ was 0.34, resulting in the prior illustrated by Fig. 2a:

$$(2.8) \quad \log \pi(\varphi_p) \sim \mathcal{N} \left(\log(16.47) - \frac{0.34^2}{2}, 0.34^2 \right)$$

This agrees well with the theoretical value of 5 to 20 cm^{-1} for broadening that is used in computational chemistry (Le Ru and Etchegoin, 2009, §2.2.5, p. 45). For Lorentzian peaks, the median was 25.27 and $\sigma(\log\{\varphi_p\})$ was 0.4. The resulting prior is shown in Fig. 2c. These prior distributions overlap, although the Lorentzian peaks tend towards larger scale parameters. This is consistent with the FWHM, since rescaling the prior for the squared exponential peaks by $\sqrt{2\log 2}$ results in a distribution that is very close to the prior for the Lorentzians.

Setting a general prior for the amplitudes of the peaks is more difficult, since the amplitude depends on the concentration of the dye as well as many other factors. The amplitudes for FAM in Fig. 2b are much smaller

than for the other two dyes, even though it has the same concentration. Empirical Bayes methods could be used for calibrating this prior, although we found that this was unnecessary for the experiments in Section 4. We considered two options for the form of this distribution, which both provided a reasonably good fit to the training data. A truncated normal is suitable when the number of peaks P is assumed known. This is because it allocates prior mass to very large peaks ($A_p > 20,000$) as well as smaller peaks with amplitudes near zero. A gamma distribution with shape parameter of 3 and the same mean allocates less prior mass near zero. This makes it suited for trans-dimensional algorithms such as RJ-MCMC, since it discourages the proposal of trivially small peaks and hence keeps the total number of peaks to a minimum. The mean of the amplitudes was 3709 for the squared exponential peaks. We chose to use twice the observed variance to reflect the greater uncertainty about this parameter. This resulted in the truncated normal distribution illustrated in Fig. 2b:

$$(2.9) \quad \pi(A_p) \sim \mathcal{N}(3709, 5598^2) \mathbb{I}_{(>0)}$$

For Lorentzian peaks, the mean was 3449 and twice the observed variance was 5672^2 . This prior is shown in Fig. 2d.

There are other ways to obtain prior information for these parameters, apart from the painstaking process of manual peak fitting described above. The locations and amplitudes of the peaks can be predicted from the chemical structure of the molecule using time-dependent density functional theory (TD-DFT; Van Caillie and Amos, 2000; Jensen, Aikens and Schatz, 2008). This quantum mechanical model scales geometrically with the number of atoms in the molecule, requiring a large amount of computational resources. Examples of dyes with predicted Raman spectra include rhodamine 6G (Watanabe et al., 2005), crystal violet (Kleinman et al., 2011), eosin-Y (Greeneltch et al., 2012) and II-MB-114 (Kearns et al., 2016). There are also databases available of known Raman spectra, such as RRUFF (Lafuente et al., 2015), SDBS (AIST), and the Raman Spectroscopic Library of Natural and Synthetic Pigments (Bell, Clark and Gibbs, 1998). Once our model has been fitted to a dye, the posterior distribution for the spectral signature can be used as an informative prior to analyse independent samples of the same dye. This process of learning from successive observations is known as sequential Bayesian updating (Bernardo and Smith, 1994, ch. 2).

2.2. Baseline correction. In contrast with the Raman peaks, little is known about the processes driving the baseline. The shape of the baseline can vary considerably between experiments and even occasionally between technical replicates. The main property that distinguishes the baseline

from the other components of the signal is its smoothness. For this reason, we have chosen to model the baseline function as a penalised B-spline (Eilers and Marx, 1996):

$$(2.10) \quad \xi_i(\tilde{\nu}) = \sum_{m=1}^M B_m(\tilde{\nu}) \alpha_{i,m}$$

where B_m are the basis functions, M is the total number of splines, and $\alpha_{i,m}$ are the coefficients of the baseline for the i th observation. We use equally spaced knots 10 cm^{-1} apart, so that M is typically ≈ 200 . If the choice of knot locations is a concern, then a smoothing spline (Eubank, 1999) could be used instead. Razul, Fitzgerald and Andrieu (2003) used an RJ-MCMC algorithm to determine the number and placement of the knots in the baseline function. As with many Bayesian regression models, (2.10) can be interpreted as a type of Gaussian process (GP; Rasmussen and Williams, 2006, §6.3). An advantage of our approach is that we employ a reduced-rank representation of the baseline function. The computational cost of estimating the spline parameters $\alpha_{i,m}$ using sparse matrix algebra is $\mathcal{O}(n)$ (Green and Silverman, 1994). This is far more scalable than the usual GP methods, which require $\mathcal{O}(n^3)$ operations to invert the covariance matrix. Alternative methods for fast GP fitting include the fixed-rank kriging of Cressie and Johannesson (2008), the Markov random field representation of Lindgren, Rue and Lindström (2011), and the nearest-neighbour GP of Datta et al. (2016).

Given the current estimate of the spectral signature $\hat{s}(\tilde{\nu})$, new values for the spline coefficients can be proposed from the conditional posterior (Ruppert, Wand and Carroll, 2003, ch. 16), as shown in Appendix A.1. The joint prior on the spline coefficients is multivariate normal:

$$(2.11) \quad \pi(\alpha_{i,\cdot}) \sim \mathcal{N}_M(0, \lambda^{-6} \mathbf{D}^{-1})$$

where \mathbf{D} is a diagonal $M \times M$ matrix and λ is a hyperparameter that controls the degree of smoothing.

3. Bayesian Computation. The amplitudes of the Raman peaks are completely dependent on the position of the baseline, and vice-versa. This suggests that the best way to perform peak fitting and baseline correction is to estimate them jointly, for example by incorporating both tasks into an iterative algorithm. However, the high degree of correlation between the position of the baseline and the amplitudes of the peaks can make many algorithms prone to becoming stuck on local modes. To address this issue, we

Algorithm 1 Marginal Metropolis-Hastings

```

1: Initialise  $\mathbf{A}^{(0)}, \boldsymbol{\varphi}^{(0)}, \boldsymbol{\alpha}^{(0)}, \sigma_\epsilon^{(0)}$ 
2: for  $t = 1, \dots, T$  do
3:   Draw random walk proposals for the peaks:  $\mathbf{A}', \boldsymbol{\varphi}'$ 
4:   Propose baseline  $\alpha'_{i,\cdot} \sim q(\alpha_{i,\cdot} \mid \mathbf{Y}_{1:n_y, \tilde{\nu}}, \mathbf{A}', \boldsymbol{\varphi}', \sigma_\epsilon^{(t-1)})$ 
5:   Propose  $\sigma'_\epsilon \sim q(\sigma_\epsilon \mid \mathbf{Y}_{1:n_y, \tilde{\nu}}, \mathbf{A}', \boldsymbol{\varphi}', \alpha'_{i,\cdot})$ 
6:   Compute the marginal MH acceptance ratio,  $\rho_t$ 
7:   Draw  $u \sim \text{Uniform}[0, 1]$ 
8:   if  $u < \min(1, \rho_t)$  then
9:      $\{\mathbf{A}^{(t)}, \boldsymbol{\varphi}^{(t)}, \alpha_{i,\cdot}^{(t)}, \sigma_\epsilon^{(t)}\} \leftarrow \{\mathbf{A}', \boldsymbol{\varphi}', \alpha'_{i,\cdot}, \sigma'_\epsilon\}$ 
10:  else
11:     $\{\mathbf{A}^{(t)}, \boldsymbol{\varphi}^{(t)}, \alpha_{i,\cdot}^{(t)}, \sigma_\epsilon^{(t)}\} \leftarrow \{\mathbf{A}^{(t-1)}, \boldsymbol{\varphi}^{(t-1)}, \alpha_{i,\cdot}^{(t-1)}, \sigma_\epsilon^{(t-1)}\}$ 
12:  end if
13: end for

```

adopt a marginal Metropolis-Hastings (MH) algorithm, as shown in Algorithm 1. This algorithm targets the joint posterior of the model parameters, $\pi(\mathbf{A}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon \mid \mathbf{Y}_{1:n_y, \tilde{\nu}})$. In this version of the algorithm, the number P and locations ℓ_p of the peaks are assumed known. Since we use conjugate priors, the full conditional distributions for the baseline coefficients $\boldsymbol{\alpha}$ and the additive noise σ_ϵ are available in closed form. Given random walk proposals for the parameters of the peaks, we can sample from these conditional distributions as shown in Appendices A.1 and A.2. According to the marginalisation identity of Chib (1995), the resulting MH acceptance ratio is equivalent to the posterior distribution of the peaks, with the baselines and the noise marginalised out:

$$\begin{aligned}
\rho_t &= \frac{p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \mathbf{A}', \boldsymbol{\varphi}', \boldsymbol{\alpha}', \sigma'_\epsilon) q(\boldsymbol{\alpha}^\circ \mid \cdot) q(\sigma_\epsilon^\circ \mid \cdot) \pi(\mathbf{A}') \pi(\boldsymbol{\varphi}') \pi(\boldsymbol{\alpha}') \pi(\sigma'_\epsilon)}{p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \mathbf{A}^\circ, \boldsymbol{\varphi}^\circ, \boldsymbol{\alpha}^\circ, \sigma_\epsilon^\circ) q(\boldsymbol{\alpha}' \mid \cdot) q(\sigma'_\epsilon \mid \cdot) \pi(\mathbf{A}^\circ) \pi(\boldsymbol{\varphi}^\circ) \pi(\boldsymbol{\alpha}^\circ) \pi(\sigma_\epsilon^\circ)} \\
&= \frac{p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \mathbf{A}', \boldsymbol{\varphi}') \pi(\mathbf{A}') \pi(\boldsymbol{\varphi}')}{p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \mathbf{A}^\circ, \boldsymbol{\varphi}^\circ) \pi(\mathbf{A}^\circ) \pi(\boldsymbol{\varphi}^\circ)}
\end{aligned}$$

3.1. *Quantification.* A drawback of Algorithm 1 is that it can be difficult to design efficient random walk proposals for \mathbf{A}' and $\boldsymbol{\varphi}'$. Adaptive methods such as Roberts and Rosenthal (2001); Atchadé and Rosenthal (2005); Garthwaite, Fan and Sisson (2016) can be used, but the posterior distribution contracts very quickly as the number of observations n_y and wavenumbers $n_{\tilde{\nu}} = |\tilde{\nu}|$ increases. This becomes a problem for large datasets such as in Sect. 4.2, where $n_y = 315$ and $n_{\tilde{\nu}} = 3501$. An alternative is to embed Algorithm 1 within a sequential Monte Carlo (SMC) algorithm that updates the posterior distribution using one observation at a time (Chopin, 2002).

Algorithm 2 targets a sequence of partial posteriors:

$$(3.1) \quad \pi_i(\boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon \mid \mathbf{Y}_{1:i, \tilde{\nu}}) \propto p(\mathbf{Y}_{1:i, \tilde{\nu}} \mid \boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon) \pi(\boldsymbol{\beta}) \pi(\boldsymbol{\varphi}) \pi(\boldsymbol{\alpha}) \pi(\sigma_\epsilon)$$

for $i = 0, \dots, n_y$, where π_0 is the joint prior distribution. The target distributions of the parameters $\boldsymbol{\beta}$, $\boldsymbol{\varphi}$, $\boldsymbol{\alpha}$ and σ_ϵ are represented using a collection of Q weighted parameter values, known as particles. The algorithm has four major stages: initialisation, adaptation, resampling, and mutation.

Algorithm 2 SMC

- 1: Initialise $\boldsymbol{\beta}^{(q)}, \boldsymbol{\varphi}^{(q)}, \boldsymbol{\alpha}^{(q)}, \sigma_\epsilon^{(q)} \forall q \in \{1, \dots, Q\}$
 - 2: Initialise importance weights, $w_0^{(q)} = \frac{1}{Q}$
 - 3: **for all** observations $i = 1, \dots, n$ **do**
 - 4: Update importance weights according to (3.3)
 - 5: Resample particles **if** ESS_i is below threshold
 - 6: **for all** particles $q \in \{1, \dots, Q\}$ **do**
 - 7: Update $\boldsymbol{\varphi}^{(q)}, \boldsymbol{\beta}^{(q)}, \boldsymbol{\alpha}^{(q)}, \sigma_\epsilon^{(q)}$ using Algorithm 1
 - 8: **end for**
 - 9: **end for**
-

At the initialisation stage (lines 1 and 2 of Algorithm 2) the parameter values are drawn independently from their respective prior distributions (2.8), (2.9), (2.11), and (A.5). All of the particles are assigned equal importance weights, $w_0^{(q)} = \frac{1}{Q}$. The importance weights are updated at each SMC iteration according to:

$$(3.2) \quad w_i^{(q)} \propto w_{i-1}^{(q)} \frac{p(\mathbf{Y}_{1:i, \tilde{\nu}} \mid \boldsymbol{\varphi}^{(q)}, \boldsymbol{\beta}^{(q)}, \boldsymbol{\alpha}_{1:i, \cdot}^{(q)}, \sigma_\epsilon^{(q)})}{p(\mathbf{Y}_{1:(i-1), \tilde{\nu}} \mid \boldsymbol{\varphi}^{(q)}, \boldsymbol{\beta}^{(q)}, \boldsymbol{\alpha}_{1:(i-1), \cdot}^{(q)}, \sigma_\epsilon^{(q)})}$$

$$(3.3) \quad \propto w_{i-1}^{(q)} \prod_{j=1}^{n_{\tilde{\nu}}} p(y_{i,j} \mid \boldsymbol{\varphi}^{(q)}, \boldsymbol{\beta}^{(q)}, \boldsymbol{\alpha}_{i, \cdot}^{(q)}, \sigma_\epsilon^{(q)})$$

where $p(\mathbf{Y}_{1:i, \tilde{\nu}} \mid \boldsymbol{\varphi}^{(q)}, \boldsymbol{\beta}^{(q)}, \boldsymbol{\alpha}_{1:i, \cdot}^{(q)}, \sigma_\epsilon^{(q)}) = 1$ when $i = 0$. These weights are normalised so that $\sum_{q=1}^Q w_t^{(q)} = 1$.

The weights gradually degenerate over successive iterations, which is measured by the effective sample size (ESS; Liu, 2001, pp. 34–36):

$$(3.4) \quad ESS_t = \frac{1}{\sum_{q=1}^Q (w_t^{(q)})^2}$$

At the initialisation stage, $ESS_0 = Q$. The observations $\mathbf{Y}_{1:n_y, \tilde{\nu}}$ of (2.1) are exchangeable, unlike in time series models. It can be beneficial to process the

observations in a random order, as recommended by [Chopin \(2002\)](#), to avoid having too many observations in a row with the same concentration, c_i . If desired, the observations can be grouped into batches $b_1 \subset b_2 \subset \dots \subset \mathbf{Y}_{1:n_y, \tilde{\nu}}$. Very high resolution spectra, such as when $n_{\tilde{\nu}} = 3501$, can result in rapid degeneration of the importance sampling distribution. In this case, it can be beneficial to divide a single observation into multiple batches according to a target $\zeta \in (0, 1)$ for the rate of reduction in the ESS at each iteration. This target can be tuned adaptively according to the MH acceptance rate.

When the ESS falls below a given threshold (usually set at $Q/2$), the particles are resampled according to the multinomial distribution defined by their weights. We use residual resampling ([Liu and Chen, 1998](#)), since this reduces the variance in comparison to simple multinomial draws ([Douc, Cappé and Moulines, 2005](#)). By reordering the ancestry vector that identifies which particles to resample, this operation can be performed in parallel ([Murray, Lee and Jacob, 2016](#)). After resampling, the importance weights are reset to $1/Q$ and hence $ESS_t = Q$.

The resampling step introduces duplicates into the population of particles. To reduce this redundancy, we update the parameter values using a Markov chain Monte Carlo (MCMC) kernel with invariant target distribution $\pi_i(\boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon \mid \mathbf{Y}_{1:i, \tilde{\nu}})$. One or more iterations of Algorithm 1 can be used, depending on the MH acceptance rate.

3.2. Model choice. Algorithm 2 provides a consistent, unbiased estimate of the model evidence, also known as the marginal likelihood or normalising constant ([Del Moral, Doucet and Jasra, 2006](#); [Pitt et al., 2012](#)):

$$(3.5) \quad \frac{\widehat{\mathcal{Z}}_i}{\mathcal{Z}_{i-1}} = \sum_{q=1}^Q w_{i-1}^{(q)} \prod_{j=1}^{n_{\tilde{\nu}}} p(y_{i,j} \mid \boldsymbol{\Theta}_k^{(q)})$$

$$(3.6) \quad p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid M_k) \approx \widehat{\mathcal{Z}}_n = \prod_{i=1}^{n_y} \frac{\widehat{\mathcal{Z}}_i}{\mathcal{Z}_{i-1}}$$

where $\boldsymbol{\Theta}_k^{(q)} = \{\boldsymbol{\beta}_k^{(q)}, \boldsymbol{\varphi}_k^{(q)}, \boldsymbol{\alpha}_k^{(q)}, \sigma_k^{(q)}\}$ are the parameters of model M_k for particle q and $\mathcal{Z}_0 = 1$.

The evidence for two competing models can be compared using a Bayes factor:

$$(3.7) \quad BF_{G:L} = \frac{p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid M_G)}{p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid M_L)}$$

$$(3.8) \quad = \frac{\int p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \boldsymbol{\Theta}_G, M_G) \pi(\boldsymbol{\Theta}_G \mid M_G) d\boldsymbol{\Theta}_G}{\int p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \boldsymbol{\Theta}_L, M_L) \pi(\boldsymbol{\Theta}_L \mid M_L) d\boldsymbol{\Theta}_L}$$

TABLE 1
Model evidence for the simulation study, averaged over 20 replicates

Synthetic Data	M_k	$\ln p(\mathbf{Y}_{1,\bar{\nu}} \mid M_k)$	Std. Dev.
sq. exp.	sq. exp.	-5331	83
sq. exp.	Lorentzian	-5542	94
Lorentzian	sq. exp.	-5453	72
Lorentzian	Lorentzian	-5293	88

In our case, the models under consideration differ only by their radial basis functions: M_G represents squared exponential peaks and M_L is the Lorentzian model. We have derived informative priors for these models, $\pi(\Theta_G \mid M_G)$ and $\pi(\Theta_L \mid M_L)$, using manual baseline correction and peak fitting of the same reference spectra. These priors are illustrated in Fig. 2. The algorithms for estimating the posterior distributions and the model evidence are identical, whichever function is used. In Section 4.1, we will show that only a single observation with $n_{\bar{\nu}} = 726$ is sufficient to distinguish between M_G and M_L .

4. Results.

4.1. *Model choice.* We used a simulation study to verify that the posterior distributions obtained from Algorithms 1 and 2 overlap with the true parameter values in synthetic data. The results of this study are provided in Supplement B. Under model misspecification, the posteriors for the amplitudes were underestimated, relative to the true values. Therefore, model selection is an important issue for quantification of SERS spectra. We also computed the model evidence for 20 datasets, both when the generative model for the synthetic data was identical to the model used by the algorithm, as well as under model misspecification. The results are shown in Table 1.

In both experiments, the true model had the larger model evidence in the majority of cases. When the synthetic data was generated using the squared exponential function, the natural logarithm of the Bayes factor ($\ln BF_{G:L}$) was 211 on average. According to Kass and Raftery (1995), a Bayes factor greater than 150 ($2 \ln BF_{G:L} > 10$) indicates very strong evidence in favour of M_G . Conversely, a Bayes factor less than -150 would be strong evidence in favour of M_L . However, we found that the variance of the model evidence (and hence the Bayes factor) was much larger than those guidelines would suggest. The empirical standard deviation of the log marginal over 20 replicates was 83 under the true model and 94 under model misspecification. This

TABLE 2
Model evidence for SERS spectra at two different resolutions ($n_{\bar{\nu}} = 417$ and 726) for 5 observations of 3 different dyes (TAMRA, FAM and Cy3).

Observed Data	$n_{\bar{\nu}}$	M_k	$\ln p(\mathbf{Y}_{1,\bar{\nu}} M_k)$	Std. Dev.
TAMRA	417	sq. exp.	-2796	27
		Lorentzian	-2716	52
	726	sq. exp.	-5363	13
		Lorentzian	-5317	13
FAM	417	sq. exp.	-2425	84
		Lorentzian	-2415	72
	726	sq. exp.	-4875	56
		Lorentzian	-4801	56
Cy3	417	sq. exp.	-2579	60
		Lorentzian	-2548	69
	726	sq. exp.	-5715	24
		Lorentzian	-5629	30

variation is due to a combination of Monte Carlo error as well as differences between the simulated datasets.

When the synthetic data was generated using the Lorentzian function, the mean $\ln BF_{G:L}$ was -160, again favouring the true model. However, the log Bayes factor was positive for one of the datasets. Because of this issue, we prefer to compute the model evidence for multiple datasets, rather than relying on the magnitude of the Bayes factor alone to determine the strength of the evidence.

Having established that the Bayes factor is capable of distinguishing between the two candidate kernel functions, we also applied these models to observed Raman spectra. The results of these experiments are shown in Table 2. We calculated Bayes factors for 5 observations of 3 different dyes: TAMRA, FAM and Cy3. The datasets with $n_{\bar{\nu}} = 417$ had a resolution of approximately 3 cm^{-1} . These spectra were previously analysed by [Faulds et al. \(2008\)](#) and [Zhong et al. \(2011\)](#). The spectra with $n_{\bar{\nu}} = 726$ had a resolution of 2 cm^{-1} and were analysed by [Gracie et al. \(2014\)](#). For all three dyes, the Bayes factor favoured Lorentzian peaks for data at 2 cm^{-1} resolution. However, with the coarser resolution spectra the results for Cy3 and FAM were inconclusive. The variance of the model evidence also tended to be larger for the lower resolution spectra, reflecting greater posterior uncertainty in the parameters.


4.2. *Quantification.* [Gracie et al. \(2014\)](#) performed a dilution study to estimate the LOD of T20 (TAMRA+DNA) with Ag NP, using a 100 mW laser at 532 nm. The final target concentration of the dye ranged from 0.13

TABLE 3
95% highest posterior density (HPD) intervals for the T20 dilution study.


ℓ_p (cm ⁻¹)	β_p (nM ⁻¹)	FWHM (cm ⁻¹)	LOD (nM)
460	[6.73; 17.23]	[0.00; 14.43]	[0.211; 0.784]
505	[167.95; 177.83]	[14.36; 15.52]	[0.019; 0.047]
632	[253.65; 263.16]	[11.54; 12.13]	[0.012; 0.032]
725	[19.63; 29.52]	[9.97; 16.43]	[0.123; 0.316]
752	[44.74; 54.81]	[19.49; 23.45]	[0.059; 0.156]
843	[23.48; 33.73]	[15.97; 22.26]	[0.106; 0.297]
965	[17.78; 28.09]	[12.07; 20.32]	[0.135; 0.355]
1140	[25.49; 36.11]	[20.41; 27.92]	[0.089; 0.253]
1190	[18.67; 28.99]	[11.27; 19.36]	[0.110; 0.352]
1220	[147.84; 158.30]	[17.68; 19.20]	[0.020; 0.051]
1290	[31.19; 42.49]	[15.72; 25.80]	[0.080; 0.213]
1358	[210.46; 221.96]	[17.63; 18.97]	[0.015; 0.035]
1422	[53.13; 63.99]	[15.34; 18.16]	[0.059; 0.135]
1455	[39.05; 53.87]	[20.61; 41.11]	[0.069; 0.175]
1512	[146.07; 157.28]	[20.81; 22.38]	[0.019; 0.049]
1536	[209.18; 221.03]	[14.43; 15.80]	[0.014; 0.036]
1570	[38.54; 53.67]	[23.87; 51.02]	[0.073; 0.173]
1655	[467.58; 477.91]	[17.40; 17.92]	[0.006; 0.016]

to 24.7 nM, with 21 different values of c_i . There were 3 repeats of 5 technical replicates, giving a total sample size of 315 spectra. The resolution of the spectrometer provided 3,501 wavenumbers from 250 to 2000 cm⁻¹. All of the observations were discretised with the same x axis. We truncated the spectra below 340 cm⁻¹, since there was too much noise at small Raman offsets $\Delta\tilde{\nu}$ for accurate quantification. 95% HPD intervals for the regression coefficients β_p , the FWHM and the LOD of each peak are shown in Table 3.

To verify the HPD intervals for the LOD, we can closely examine the spectra at the two lowest concentrations, 0.13 and 0.65 nM. The lower bounds for detectability of the peaks at 460 and 965 cm⁻¹ are greater than 0.13 nM, so we would not expect those peaks to be visible at that concentration. Conversely, the upper bounds for 7 of the peaks are lower than 0.13 nM, so we would expect all of those peaks to be clearly visible, as shown by Fig. 3a. There is also an eighth peak at 1455 cm⁻¹ that has been underestimated by the model. The upper bounds for all of the peaks except at 460 cm⁻¹ are lower than 0.65 nM, so at least 17 out of 18 peaks should be visible at that concentration, as shown in Fig 3b. Care must be taken when extrapolating beyond the range of the data, but we predict that overall the LOD for TAMRA is between 6 and 16 picomolar (pM).



(a) TAMRA at 0.13 nM.



(b) TAMRA at 0.65 nM.

Fig 3: Observed spectra and model fit at very low dye concentrations.

5. Discussion. Our model has enabled us to directly estimate quantities of scientific interest, such as the amplitudes, LOD and FWHM of the SERS peaks. This represents an important tool for analysing experimental data, with major advantages over the existing methods for baseline correction and quantification. Our R package (included in [Supplement A](#)) can be run directly on data from the spectrometer, without any preprocessing steps required. We provide a continuous representation of the discretised spectra, which enables comparison of datasets at different sampling resolutions. The baselines and peaks are estimated jointly, avoiding local maxima in the multimodal posterior. We have incorporated prior information from independent datasets and shown how the posterior can be incrementally updated using a sequential Monte Carlo algorithm, as more data becomes available. Bayesian credible intervals for the parameters also provide accurate estimates of the uncertainty.

This is the first study to employ Bayes factors to determine the spectral line shape of experimental data. We have compared two competing models for the broadening of the peaks, squared exponential (Gaussian) and Lorentzian (Cauchy) functions. We have demonstrated that these functions give significantly different estimates of the amplitudes, thus it is important to select the model that better represents the observed data. By computing the model evidence, these models can be compared using only a single observed spectrum. Our experiments have provided strong evidence in favour of Lorentzian peaks, which demonstrates the presence of long-range dependence in the SERS spectrum. This finding further reinforces the need for joint models to avoid bias in estimating the LOD for a dye using a single peak in isolation.

In future work, we intend to apply these methods to other Raman-active dyes to see how well they generalise. Our computational algorithms could be extended using RJ-MCMC to also estimate the number and locations of the peaks, for dyes where this information is unavailable. These methods could also be used for detection and quantification of multiplexed spectra, where several dyes are present at differing concentrations.

APPENDIX A: CONDITIONAL DISTRIBUTIONS

This appendix contains the derivation of the full posterior distribution for (2.1) and hence the conditional distributions for the model parameters α and σ_ϵ . These formulae are used in Algorithms 1 and 2 as proposal distributions for Metropolis-Hastings.

Under the assumption of constant variance σ_ϵ^2 , the joint likelihood of a

$n_y \times n_{\tilde{\nu}}$ matrix of observed data $\mathbf{Y}_{1:n_y, \tilde{\nu}}$ is *i.i.d.* Gaussian:

$$(A.1) \quad p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon) = \prod_{i=1}^{n_y} \prod_{j=1}^{n_{\tilde{\nu}}} p(y_{i,j} \mid \boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon)$$

where:

$$\begin{aligned} p(y_{i,j} \mid \boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon) &\sim \mathcal{N}(\xi_i(\nu_j) + c_i s(\nu_j), \sigma_\epsilon^2) \\ &\propto e^{-\frac{1}{2\sigma_\epsilon^2} \left(y_{i,j} - \sum_{m=1}^M B_m(\nu_j) \alpha_{i,m} - c_i \sum_{p=1}^P \beta_p \left(\frac{\varphi_p^2}{(\nu_j - \ell_p)^2 + \varphi_p^2} \right) \right)^2} \end{aligned}$$

The unnormalised posterior distribution is thus:

$$(A.2) \quad p(\boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon \mid \mathbf{Y}_{1:n_y, \tilde{\nu}}) \propto p(\mathbf{Y}_{1:n_y, \tilde{\nu}} \mid \boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \sigma_\epsilon) \pi(\boldsymbol{\beta}) \pi(\boldsymbol{\varphi}) \pi(\boldsymbol{\alpha}) \pi(\sigma_\epsilon)$$

where the priors $\pi(\varphi_p)$ are independent lognormal, as given by (2.8); $\pi(\boldsymbol{\beta})$ is multivariate normal, $\mathcal{N}_P(0, \mathbf{I}_P)$; $\pi(\boldsymbol{\alpha})$ is multivariate normal, $\mathcal{N}_M(0, \lambda^6 \mathbf{D})$; and $\pi(\sigma_\epsilon^2)$ is a conjugate, inverse gamma distribution (A.5). Note that since $A_p = c_i \beta_p$ (2.6), \mathbf{A} can be substituted for $\boldsymbol{\beta}$ in the above equations by setting $c_i = 1$.

When we update the importance weights on line 4 of Algorithm 2, we need to compute the ratio of log-likelihoods (3.3):

$$\begin{aligned} &\log \left\{ \frac{p(\mathbf{Y}_{1:i, \tilde{\nu}} \mid \boldsymbol{\varphi}, \boldsymbol{\beta}, \boldsymbol{\alpha}_{1:i, \cdot}, \sigma_\epsilon)}{p(\mathbf{Y}_{1:(i-1), \tilde{\nu}} \mid \boldsymbol{\varphi}, \boldsymbol{\beta}, \boldsymbol{\alpha}_{1:(i-1), \cdot}, \sigma_\epsilon)} \right\} \\ &= -\frac{in_{\tilde{\nu}} - (i-1)n_{\tilde{\nu}}}{2} \log\{2\pi\sigma_\epsilon^2\} - \frac{\delta_{1:i}^2}{2\sigma_\epsilon^2} + \frac{\delta_{1:(i-1)}^2}{2\sigma_\epsilon^2} \\ &= -\frac{n_{\tilde{\nu}}}{2} \log\{2\pi\sigma_\epsilon^2\} - \frac{1}{2\sigma_\epsilon^2} \sum_{j=1}^{n_{\tilde{\nu}}} (y_{i,j} - \xi_i(\nu_j) - c_i s(\nu_j))^2 \end{aligned}$$

where $\delta_{1:i}^2 = \sum_{t=1}^i \sum_{j=1}^{n_{\tilde{\nu}}} (y_{t,j} - \xi_t(\nu_j) - c_t s(\nu_j))^2$ is the sum of squared differences.

A.1. Baseline. The basis functions in (2.10) for a penalised cubic spline can be expanded as follows:

$$(A.3) \quad \xi_i(\tilde{\nu}) = \gamma_{i,0} + \tilde{\nu} \gamma_{i,1} + \tilde{\nu}^2 \gamma_{i,2} + \tilde{\nu}^3 \gamma_{i,3} + \sum_{j=1}^K \psi_{i,j}(\tilde{\nu} - \kappa_j)_+^3$$

where γ are the coefficients of the polynomial functions of the wavenumbers $\tilde{\nu}$ and ψ are the coefficients of the knots, $\kappa = \kappa_1, \dots, \kappa_K$. Thus, the vector $\alpha_{i,\cdot} = (\gamma_{i,\cdot}, \psi_{i,\cdot})$ and $M = K + 4$. Let \mathbf{B} be the $|\tilde{\mathcal{V}}| \times M$ matrix of basis functions given by (A.3) and let \mathbf{D} be a $M \times M$ diagonal matrix, $\text{diag}(0_4, 1_K)$. Given the current estimate of the spectral signature, $\hat{s}(\tilde{\nu}) = c_i \sum_{p=1}^P \beta_p \left(\frac{\varphi_p^2}{(\tilde{\nu} - \ell_p)^2 + \varphi_p^2} \right)$, the conditional posterior for $\alpha_{i,\cdot}$ is given by:

$$(A.4) \quad \pi(\alpha_{i,\cdot} \mid \mathbf{y}_i, \boldsymbol{\varphi}, \boldsymbol{\beta}, \sigma_\epsilon^2) \sim \mathcal{N}(\boldsymbol{\Sigma}_\lambda \mathbf{B}^T (\mathbf{y}_i - \hat{s}(\tilde{\nu})))^T, \sigma_\epsilon^2 \boldsymbol{\Sigma}_\lambda)$$

where $\boldsymbol{\Sigma}_\lambda = (\mathbf{B}^T \mathbf{B} + \lambda^6 \mathbf{D})^{-1}$, a sparse $M \times M$ matrix that can be precomputed for a given value of λ and set of wavenumbers $\tilde{\mathcal{V}}$. We have implemented the matrix-vector multiply in (A.4) using `RcppEigen` (Bates and Eddelbuettel, 2013).

A.2. Noise. In Algorithm 2, the variance of the additive Gaussian noise is included in the state space of the SMC particles. We assign a conjugate, inverse-Gamma prior with hyperparameters ν_0 and ζ_0 :

$$(A.5) \quad \pi\left(\frac{1}{\sigma_\epsilon^2}\right) \sim \mathcal{Ga}\left(\frac{\nu_0}{2}, \frac{2}{\nu_0 \zeta_0^2}\right)$$

The conditional posterior distribution is given by:

$$\begin{aligned} \pi(\sigma_\epsilon^2 \mid \boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\alpha}, \mathbf{Y}) &= (2\pi\sigma_\epsilon^2)^{-\frac{n|\tilde{\mathcal{V}}|}{2}} \exp\left\{-\frac{\delta^2}{2\sigma_\epsilon^2}\right\} (\sigma_\epsilon^2)^{-\frac{\nu_0}{2}-1} \exp\left\{-\frac{\nu_0 \zeta_0^2}{2\sigma_\epsilon^2}\right\} \\ &\propto (\sigma_\epsilon^2)^{-\frac{\nu_0 + n|\tilde{\mathcal{V}}|}{2}-1} \exp\left\{-\frac{\nu_0 \zeta_0^2 + \delta^2}{2\sigma_\epsilon^2}\right\} \end{aligned}$$

The hyperparameters of the conditional posterior are thus:

$$(A.6) \quad \nu' = \nu_0 + n|\tilde{\mathcal{V}}|$$

$$(A.7) \quad \zeta' = \sqrt{\frac{\nu_0 \zeta_0^2 + \delta^2}{\nu'}}$$

ACKNOWLEDGEMENTS

We warmly thank Daniel Simpson, Kerrie Mengersen, Christian Robert, Hayleigh May, Ivan Ramos Sasselli, and Steve Asiala for insightful discussions. This work was funded by the UK EPSRC programme grant, “*In Situ Nanoparticle Assemblies for Healthcare Diagnostics and Therapy*” (ref: EP/L014165/1) and an Award for Postdoctoral Collaboration from the EPSRC Network on Computational Statistics & Machine Learning (ref: EP/K009788/2).

SUPPLEMENTARY MATERIAL

Supplement A: R package serrsBayes

([http://lib.stat.cmu.edu/aoas/???/???;](http://lib.stat.cmu.edu/aoas/???/???) .tar.gz). An R source package (Gnu zipped tarball) containing R and C++ code for peak fitting and baseline correction of Raman spectra. These functions implement the sequential Monte Carlo and Metropolis-Hastings algorithms that we have introduced in this article.

Supplement B: Additional results

([http://lib.stat.cmu.edu/aoas/???/???;](http://lib.stat.cmu.edu/aoas/???/???) .pdf). Posterior distributions for the peaks and baselines of 3 Raman-active dyes (FAM, TAMRA, and Cy3) at 2 different resolutions ($n_{\tilde{\nu}} = 417$ and 726). We also provide additional results for the simulation study of Section 4.1 and the concentration study of Section 4.2.

REFERENCES

- AIST Spectral Database for Organic Compounds (SDBS). National Institute of Advanced Industrial Science and Technology, Tokyo, Japan.
- ATCHADÉ, Y. F. and ROSENTHAL, J. S. (2005). On adaptive Markov chain Monte Carlo algorithms. *Bernoulli* **11** 815–828.
- BATES, D. and EDELBUETTTEL, D. (2013). Fast and Elegant Numerical Linear Algebra Using the RcppEigen Package. *J. Stat. Soft.* **52** 1–24.
- BELL, I. M., CLARK, R. J. H. and GIBBS, P. J. (1998). Raman Spectroscopic Library of Natural and Synthetic Pigments. University College London, UK.
- BERNARDO, J. M. and SMITH, A. F. M. (1994). *Bayesian Theory*. John Wiley & Sons, Chichester, UK.
- BOELEN, H. F. M., EILERS, P. H. C. and HANKEMEIER, T. (2005). Sign Constraints Improve the Detection of Differences between Complex Spectral Data Sets: LC-IR As an Example. *Anal. Chem.* **77** 7998–8007.
- BUTLER, H. J., ASHTON, L., BIRD, B., CINQUE, G., CURTIS, K., DORNEY, J., ESMONDE-WHITE, K., FULLWOOD, N. J., GARDNER, B., MARTIN-HIRSCH, P. L., WALSH, M. J., MCAINSH, M. R., STONE, N. and MARTIN, F. L. (2016). Using Raman spectroscopy to characterize biological materials. *Nat. Protocols* **11** 664–687.
- CAI, T. T., ZHANG, D. and BEN-AMOTZ, D. (2001). Enhanced chemical classification of Raman images using multiresolution wavelet transformation. *Appl. Spectrosc.* **55** 1124–30.
- CHIB, S. (1993). Bayes regression with autoregressive errors: A Gibbs sampling approach. *J. Econometrics* **58** 275–294.
- CHIB, S. (1995). Marginal Likelihood from the Gibbs Output. *J. Am. Stat. Assoc.* **90** 1313–1321.
- CHOPIN, N. (2002). A Sequential Particle Filter Method for Static Models. *Biometrika* **89** 539–551.
- CRESSIE, N. and JOHANNESSON, G. (2008). Fixed rank kriging for very large spatial data sets. *J. R. Stat. Soc. Ser. B* **70** 209–226.

- DATTA, A., BANERJEE, S., FINLEY, A. O. and GELFAND, A. E. (2016). Hierarchical Nearest-Neighbor Gaussian Process Models for Large Geostatistical Datasets. *J. Am. Stat. Assoc.* In press.
- DEL MORAL, P., DOUCET, A. and JASRA, A. (2006). Sequential Monte Carlo samplers. *J. R. Stat. Soc. Ser. B* **68** 411–436.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B* **39** 1–38.
- DOUC, R., CAPPÉ, O. and MOULINES, É. (2005). Comparison of resampling schemes for particle filtering. In *Proc. 4th Int. Symp. Image and Signal Processing and Analysis (ISPA)* 64–69. IEEE, Zagreb, Croatia.
- EILERS, P. H. C. and MARX, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statist. Sci.* **11** 89–121.
- EUBANK, R. L. (1999). *Nonparametric Regression and Spline Smoothing*, 2nd ed. *Statistics: textbooks and monographs* **157**. Marcel Dekker, New York, NY.
- FAULDS, K., JARVIS, R., SMITH, W. E., GRAHAM, D. and GOODACRE, R. (2008). Multiplexed detection of six labelled oligonucleotides using surface enhanced resonance Raman scattering (SERRS). *Analyst* **133** 1505–12.
- GALLOWAY, C. M., LE RU, E. C. and ETCHEGOIN, P. G. (2009). An Iterative Algorithm for Background Removal in Spectroscopy by Wavelet Transforms. *Appl. Spectrosc.* **63** 1370–1376.
- GARTHWAITE, P. H., FAN, Y. and SISSON, S. A. (2016). Adaptive optimal scaling of Metropolis-Hastings algorithms using the Robbins-Monro process. *Commun. Stat. Theory*. In press.
- GOLDBERG, H. S. and PERSHAN, P. S. (1973). Raman line shapes in liquid CH₃I and CD₃I. *J. Chem. Phys.* **58** 3816–3827.
- GRACIE, K., CORREA, E., MABBOTT, S., DOUGAN, J. A., GRAHAM, D., GOODACRE, R. and FAULDS, K. (2014). Simultaneous detection and quantification of three bacterial meningitis pathogens by SERS. *Chem. Sci.* **5** 1030–40.
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82** 711–732.
- GREEN, P. J. and SILVERMAN, B. W. (1994). *Nonparametric regression and generalized linear models: a roughness penalty approach*. *Monographs on Statistics and Applied Probability* **58**. Chapman & Hall/CRC Press, Boca Raton, FL.
- GREENELTCH, N. G., DAVIS, A. S., VALLEY, N. A., CASADIO, F., SCHATZ, G. C., VAN DUYN, R. P. and SHAH, N. C. (2012). Near-Infrared Surface-Enhanced Raman Spectroscopy (NIR-SERS) for the Identification of Eosin Y: Theoretical Calculations and Evaluation of Two Different Nanoplasmonic Substrates. *J. Phys. Chem. A* **116** 11863–11869.
- HE, S., ZHANG, W., LIU, L., HUANG, Y., HE, J., XIE, W., WU, P. and DU, C. (2014). Baseline correction for Raman spectra using an improved asymmetric least squares method. *Anal. Methods* **6** 4402–4407.
- JENSEN, L., AIKENS, C. M. and SCHATZ, G. C. (2008). Electronic structure methods for studying surface-enhanced Raman scattering. *Chem. Soc. Rev.* **37** 1061–1073.
- JONES, J. C., MCLAUGHLIN, C., LITTLEJOHN, D., SADLER, D. A., GRAHAM, D. and SMITH, W. E. (1999). Quantitative Assessment of Surface-Enhanced Resonance Raman Scattering for the Analysis of Dyes on Colloidal Silver. *Anal. Chem.* **71** 596–601.
- KASS, R. E. and RAFTERY, A. E. (1995). Bayes Factors. *J. Am. Stat. Assoc.* **90** 773–795.
- KEARNS, H., SENGUPTA, S., RAMOS SASSELLI, I., BROMLEY III, L., FAULDS, K., TUTTLE, T., BEDICS, M. A., DETTY, M. R., VELARDE, L., GRAHAM, D. and SMITH, W. E. (2016). Elucidation of the bonding of a near infrared dye to hollow gold nanospheres –

- a chalcogen tripod. *Chem. Sci.* In press.
- KIM, S., OUYANG, M., JEONG, J., SHEN, C. and ZHANG, X. (2014). A new method of peak detection for analysis of comprehensive two-dimensional gas chromatography mass spectrometry data. *Ann. Appl. Stat.* **8** 1209–1231.
- KLEINMAN, S. L., RINGE, E., VALLEY, N., WUSTHOLZ, K. L., PHILLIPS, E., SCHEIDT, K. A., SCHATZ, G. C. and DUYNE, R. P. V. (2011). Single-Molecule Surface-Enhanced Raman Spectroscopy of Crystal Violet Isotopologues: Theory and Experiment. *J. Am. Chem. Soc.* **133** 4115–4122.
- LAFUENTE, B., DOWNS, R. T., YANG, H. and STONE, N. (2015). The power of databases: the RRUFF project. In *Highlights in Mineralogical Crystallography* (T. Armbruster and R. M. Danisi, eds.) 1–30. W. De Gruyter, Berlin, Germany.
- LE RU, E. C. and ETCHGOIN, P. G. (2009). *Principles of Surface-Enhanced Raman Spectroscopy and Related Plasmonic Effects*. Elsevier, Amsterdam, Netherlands.
- LIEBER, C. A. and MAHADEVAN-JANSEN, A. (2003). Automated Method for Subtraction of Fluorescence from Biological Raman Spectra. *Appl. Spectrosc.* **57** 1363–1367.
- LILAND, K. H., ALMØY, T. and MEVIK, B.-H. (2010). Optimal Choice of Baseline Correction for Multivariate Calibration of Spectra. *Appl. Spectrosc.* **64** 234A–268A.
- LINDGREN, F., RUE, H. and LINDSTRÖM, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B* **73** 423–498.
- LIU, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. Springer, New York, NY.
- LIU, J. S. and CHEN, R. (1998). Sequential Monte Carlo Methods for Dynamic Systems. *J. Am. Stat. Assoc.* **93** 1032–1044.
- McLACHLAN, G. J. and KRISHNAN, T. (2008). *The EM algorithm and extensions*, 2nd ed. John Wiley & Sons, Hoboken, NJ.
- MURRAY, L. M., LEE, A. and JACOB, P. E. (2016). Parallel resampling in the particle filter. *J. Comput. Graph. Stat.* In press.
- PITT, M. K., DOS SANTOS SILVA, R., GIORDANI, P. and KOHN, R. (2012). On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *J. Econometrics* **171** 134–151.
- RAMSAY, J. O. and SILVERMAN, B. W. (2005). *Functional Data Analysis*, 2nd ed. Springer, New York, NY.
- RASMUSSEN, C. E. and WILLIAMS, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA.
- RAZUL, S. G., FITZGERALD, W. J. and ANDRIEU, C. (2003). Bayesian model selection and parameter estimation of nuclear emission spectra using RJMCMC. *Nucl. Instrum. Meth. A* **497** 492–510.
- ROBERTS, G. O. and ROSENTHAL, J. S. (2001). Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.* **16** 351–367.
- DE ROOI, J. J. and EILERS, P. H. C. (2012). Mixture models for baseline estimation. *Chemometr. Intell. Lab.* **117** 56–60.
- RUCKSTUHL, A. F., JACOBSON, M. P., FIELD, R. W. and DODD, J. A. (2001). Baseline subtraction using robust local regression estimation. *J. Quant. Spectrosc. Ra.* **68** 179–193.
- RUPPERT, D. R., WAND, M. P. and CARROLL, R. J. (2003). *Semiparametric Regression*. Cambridge University Press, Cambridge, UK.
- SPANIER, J. E., ROBINSON, R. D., ZHANG, F., CHAN, S.-W. and HERMAN, I. P. (2001). Size-dependent properties of CeO₂-y nanoparticles as studied by Raman scattering. *Phys. Rev. B* **64** 245407.

- VAN CAILLIE, C. and AMOS, R. D. (2000). Raman intensities using time dependent density functional theory. *Phys. Chem. Chem. Phys.* **2** 2123–29.
- VARMUZA, K. and FILZMOSER, P. (2009). *Introduction to Multivariate Statistical Analysis in Chemometrics*. Chapman & Hall/CRC Press, Boca Raton, FL.
- WANG, Y., ZHOU, X., WANG, H., LI, K., YAO, L. and WONG, S. T. C. (2008). Reversible jump MCMC approach for peak identification for stroke SELDI mass spectrometry using mixture model. *Bioinformatics* **24** i407–i413.
- WATANABE, H., HAYAZAWA, N., INOUE, Y. and KAWATA, S. (2005). DFT Vibrational Calculations of Rhodamine 6G Adsorbed on Silver: Analysis of Tip-Enhanced Raman Spectroscopy. *J. Phys. Chem. B* **109** 5012–5020.
- ZHONG, M., GIROLAMI, M., FAULDS, K. and GRAHAM, D. (2011). Bayesian methods to detect dye-labelled DNA oligonucleotides in multiplexed Raman spectra. *J. R. Stat. Soc. Ser. C* **60** 187–206.
- ZHOU, H. and LANGE, K. L. (2010). On the Bumpy Road to the Dominant Mode. *Scand. J. Statist.* **37** 612–631.

DEPARTMENT OF STATISTICS
UNIVERSITY OF WARWICK
COVENTRY, CV4 7AL, UK
E-MAIL: M.T.Moores@warwick.ac.uk
Jake.Carson@warwick.ac.uk
M.Girolami@warwick.ac.uk
URL: <http://www.warwick.ac.uk/mmoores>

CENTRE FOR MOLECULAR NANOMETROLOGY
DEPARTMENT OF PURE & APPLIED CHEMISTRY
UNIVERSITY OF STRATHCLYDE
TECHNOLOGY AND INNOVATION CENTRE
99 GEORGE STREET, GLASGOW, G1 1RD, UK
E-MAIL: duncan.graham@strath.ac.uk
karen.faulds@strath.ac.uk
kirsten.gracie@strath.ac.uk
URL: <http://www.strath.ac.uk/chemistry/>

THE ALAN TURING INSTITUTE FOR DATA SCIENCE
BRITISH LIBRARY
96 EUSTON ROAD, LONDON, NW1 2DB, UK
E-MAIL: mgirolami@turing.ac.uk
URL: <http://www.turing.ac.uk>